

What is claimed is:

[Claim 1] 1. A method for reliably storing data in a computer system, comprising:

receiving a piece of data to be stored at a storage system;
writing a first copy of the data to the storage system according to a first data redundancy scheme;
writing a second copy of the data to the storage system according to a second data redundancy scheme;
maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

[Claim 2] 2. The method of claim 1 wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of blocks of user data.

[Claim 3] 3 The method of claim 2 wherein the metadata is maintained by using a skeleton file system that replicates all the namespace and attribute information of the stored data and a set of delta files each of which records all the data updates that have been made to a corresponding file and that have not been copied to the backup storage system.

[Claim 4] 4. The method of claim 1 wherein the maintained metadata comprises address information associated with where the data is stored in the storage system according to the first data redundancy scheme.

[Claim 5] 5. The method of claim 1 further comprises recognizing a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme.

[Claim 6] 6. The method of claim 5 further comprises responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.

[Claim 7] 7. The method of claim 6 wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises, utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the storage system.

[Claim 8] 8. The method of claim 7 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

[Claim 9] 9. The method of claim 6 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to the second data redundancy scheme;

retrieving the data written to the storage system according to the second data redundancy scheme, where the data has not yet been copied to the backup storage system;

retrieving the data written to the backup storage system; and

rebuilding an accurate copy of the data stored according to the first data redundancy scheme, by utilizing the maintained metadata in combination with the retrieved data written to the storage system according to the second data redundancy scheme and the retrieved data written to the backup storage system.

[Claim 10] 10. The method of claim 5 wherein failure comprises a failure of a plurality of physical disks in the storage system.

[Claim 11] 11. A disaster recovery system, comprising:

a plurality of disks;

a storage controller attached to the disks, wherein the storage controller receives data to be stored on the plurality of physical disks;

first data redundancy scheme logic included within the storage controller for writing a copy of the data received by the storage controller to the physical disks according to a first data redundancy scheme;
second data redundancy scheme logic included within the storage controller for writing a second copy of the data received by the storage controller to the physical disks according to a second data redundancy scheme;
metadata of the data written to the storage system according to the second data redundancy scheme;
data backup logic to copy the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
maintenance logic to remove the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

[Claim 12] 12. The system of claim 11 wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of blocks of user data.

[Claim 13] 13. The method of claim 12 wherein the metadata is maintained by using a skeleton file system that replicates all the namespace and attribute information of the stored data and a set of delta files each of which records all the data updates that have been made to a corresponding file and that have not been copied to the backup storage system.

[Claim 14] 14. The system of claim 11 wherein the maintained metadata comprises address information associated with where the data is stored in the storage system according to the first data redundancy scheme.

[Claim 15] 15. The system of claim 11 further comprises monitoring logic to recognize a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme.

[Claim 16] 16. The system of claim 15 further comprises responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.

[Claim 17] 17. The system of claim 16 wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises, utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the storage system.

[Claim 18] 18. The system of claim 17 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

[Claim 19] 19. The system of claim 16 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to the second data redundancy scheme;

retrieving the data written to the storage system according to the second data redundancy scheme, where the data has not yet been copied to the backup storage system;

retrieving the data written to the backup storage system; and

rebuilding an accuracy copy of the data stored according to the first data redundancy scheme, by utilizing the maintained metadata in combination with the retrieved data written to the storage system according to the second data redundancy scheme and the retrieved data written to the backup storage system.

[Claim 20] 20. The system of claim 15 wherein failure comprises a failure of a plurality of physical disks in the storage system.

[Claim 21] 21 A computer program product having instruction codes for reliably storing data in a computer system, comprising:

a set of instruction codes receiving a piece of data to be stored at a storage system;

a set of instruction codes for writing a first copy of the data to the storage system according to a first data redundancy scheme;

a set of instruction codes for writing a second copy of the data to the storage system according to a second data redundancy scheme;

a set of instruction codes for maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
a set of instruction codes for copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
a set of instruction codes for removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

[Claim 22] 22. A method for deploying a disaster recovery service in a computer system, comprising:

integrating computer readable code into a system for receiving a piece of data to be stored at a storage system;
integrating computer readable code into a system for writing a first copy of the data to the storage system according to a first data redundancy scheme;
integrating computer readable code into a system for writing a second copy of the data to the storage system according to a second data redundancy scheme;
integrating computer readable code into a system for maintaining metadata of the data written to the storage system according to the second data redundancy scheme;
integrating computer readable code into a system for copying the data written to the storage system according to the second data redundancy scheme to a backup storage system, wherein the copying is performed in response to a defined condition being met; and
integrating computer readable code into a system for removing the data written to the storage system according to the second data redundancy scheme after it has been copied to the backup storage system.

[Claim 23] 23. The method of claim 22 wherein the stored data includes both blocks of user data and attributes such as file names that are associated with collections of blocks of user data.

[Claim 24] 24. The method of claim 23 wherein the metadata is maintained by using a skeleton file system that replicates all the namespace and attribute information of the stored data and a set of delta files each of which records all the data updates that have been made to a corresponding file and that have not been copied to the backup storage system.

[Claim 25] 25. The method of claim 22 wherein the maintained metadata comprises address information associated with where the data is stored in the storage system according to the first data redundancy scheme.

[Claim 26] 26. The method of claim 22 further comprises integrating computer readable code for recognizing a failure in the storage system which compromises the reliability of the data stored according to the first data redundancy scheme.

[Claim 27] 27. The method of claim 26 further comprises integrating computer readable code for responding to the failure by rebuilding an accurate copy of the data stored according to the first data redundancy scheme.

[Claim 28] 28. The method of claim 27 wherein rebuilding an accurate copy of the data stored according to the first data redundancy scheme comprises, utilizing the maintained metadata as a blueprint for combining data retrieved from the backup storage system and data remaining in the storage system.

[Claim 29] 29. The method of claim 28 wherein data remaining in the storage system comprises data stored according to the second data redundancy scheme.

[Claim 30] 30. The method of claim 27 wherein rebuilding an accurate copy of the data comprises:

retrieving the maintained metadata of the data written to the storage system according to the second data redundancy scheme;

retrieving the data written to the storage system according to the second data redundancy scheme, where the data has not yet been copied to the backup storage system;

retrieving the data written to the backup storage system; and

rebuilding an accurate copy of the data stored according to the first data redundancy scheme, by utilizing the maintained metadata in combination with the retrieved data written to the storage system according to the second data redundancy scheme and the retrieved data written to the backup storage system.

[Claim 31] 31.The method of claim 26 wherein failure comprises a failure of a plurality of physical disks in the storage system.